

Explicit Multiregister Measurements for Hidden Subgroup Problems

Cristopher Moore
University of New Mexico

Alexander Russell
University of Connecticut

Abstract

We present an explicit measurement in the Fourier basis that solves an important case of the Hidden Subgroup Problem, including the case to which Graph Isomorphism reduces. This entangled measurement uses $k = \log_2 |G|$ registers, and each of the 2^k subsets of the registers contributes some information. While this does not, in general, yield an efficient algorithm, it generalizes the relationship between Subset Sum and the HSP in the dihedral group, and sheds some light on how quantum algorithms for Graph Isomorphism might work.

1 Introduction: The Hidden Subgroup Problem

Many problems of interest in quantum computing can be expressed as, or reduced to, an instance of the *Hidden Subgroup Problem* (HSP). We are given a group G and a function f with the promise that, for some subgroup $H \subseteq G$, f is invariant precisely under translation by H : that is, $f(g_1) = f(g_2)$ if and only if $g_1 = g_2 h$ for some $h \in H$. We then wish to determine the subgroup H . Every known efficient algorithm for this problem—and, indeed, almost every quantum algorithm that provides an exponential speedup over the best known classical algorithm—uses the approach of *Fourier sampling* [5]. By preparing a uniform superposition over the elements of G , querying the function f , and then measuring the value of f , we obtain a uniform superposition over one of the (left) cosets of H ,

$$|cH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle$$

where c is a uniformly random element of G . Alternately, we can view this as a mixed state over the left cosets, the *coset state*, with density matrix

$$\rho = \frac{1}{|G|} \sum_{c \in G} |cH\rangle \langle cH| .$$

We then carry out the quantum Fourier transform on $|cH\rangle$, or equivalently ρ , and measure the result.

For example, in Simon's problem [34], $G = \mathbb{Z}_2^n$ and there is some y such that $f(x) = f(x+y)$ for all x ; in this case $H = \{0, y\}$ and we wish to identify y . In Shor's factoring algorithm [33] G is the group \mathbb{Z}_n^* where n is the number we wish to factor, $f(x) = r^x \bmod n$ for a random $r < n$, and H is the subgroup of \mathbb{Z}_n^* whose index is the multiplicative order of r . (However, since $|\mathbb{Z}_n^*|$ is unknown, we actually perform the Fourier transform over \mathbb{Z}_q for some $q = O(n^2)$; see [33] or [13, 14].) In both these algorithms, G is abelian, and it is not hard to see that for any abelian group a polynomial number¹ of experiments of this type allow us to determine H . In essence, each experiment yields a random element of the dual space H^\perp perpendicular to H 's characteristic function, and as soon as these elements span H^\perp we can determine a set of generators for H by linear algebra.

While the *nonabelian* hidden subgroup problem appears to be much more difficult, solving it would provide enormous benefits. In particular, solving the HSP for the symmetric group S_n would provide an

¹Throughout the paper, the terms polynomial, subexponential, etc. refer to a function of $\log |G|$.

efficient quantum algorithm for the Graph Automorphism and Graph Isomorphism problems (see e.g. [21] for a review). Let G_1, G_2 be two rigid, connected graphs of size n , and let $H \subset S_{2n}$ be the automorphism group of their disjoint union. If $G_1 \cong G_2$, then $H = \{1, m\}$ is of order 2, consisting of the identity and an involution m composed of n disjoint transpositions; if $G_1 \not\cong G_2$, then H is the trivial subgroup consisting only of the identity. Thus even distinguishing subgroups of order 2 from the trivial subgroup would be sufficient to solve this case of Graph Isomorphism. Other important motivations include the relationship between the HSP on the dihedral group and hidden shift problems [6] and cryptographically important cases of the Shortest Lattice Vector problem [30].

So far, explicit polynomial-time quantum algorithms for the HSP are known only for a few families of nonabelian groups [3, 10, 12, 15, 18, 20, 28, 31]. However, the basic idea of Fourier sampling can certainly be extended to the nonabelian case. Fourier basis functions are homomorphisms $\phi : G \rightarrow \mathbb{C}$ such as the familiar $\phi_k(x) = e^{2\pi i k x / n}$ when G is the cyclic group \mathbb{Z}_n . In the nonabelian case, one instead considers *representations* of G , namely homomorphisms $\sigma : G \rightarrow U(V)$ where $U(V)$ is the group of unitary matrices acting on some vector space V of dimension d_σ . The *irreducible* representations are those which are not isomorphic to direct sums of representations on lower-dimensional subspaces, and we denote the set of irreducibles by \widehat{G} . We refer the reader to [11] for an introduction. We denote the set of functions $\psi : G \rightarrow \mathbb{C}$ with $\|\psi\|^2 = 1$, i.e., the Hilbert space of a group-valued register, as $\mathbb{C}[G]$; then the quantum Fourier transform consists of transforming vectors in $\mathbb{C}[G]$ from the basis $\{|g\rangle \mid g \in G\}$ to the basis $|\sigma, i, j\rangle$ where σ is the isomorphism type, or “name,” of an irreducible representation and $1 \leq i, j \leq d_\sigma$ index a row and column (in a chosen basis for V). This transformation can be carried out efficiently for a wide variety of groups [4, 17, 27].

Several varieties of measurement in the Fourier basis have been proposed. *Weak Fourier sampling* consists of measuring just the name σ of the irreducible representation. *Strong Fourier sampling* consists of measuring the name σ and the column j in a basis of our choice. (As the state ρ is mixed uniformly over the left cosets, it is easy to show that measuring the row provides no information). As an intermediate notion, one can also consider measuring the column in a *random* basis for V .

Unfortunately, a series of negative results have shown that these types of measurement will not succeed in solving the Hidden Subgroup Problem in the cases we care most about—in particular, the case relevant to Graph Isomorphism [15, 12, 22]. In particular, Moore, Russell and Schulman [25] showed that strong Fourier sampling fails, in the sense that we need an exponential number of experiments on single coset states to distinguish the order-2 subgroups from the trivial subgroup.

However, there is still reason for hope. In the above description of Fourier sampling, f is queried just once, giving a coset state on a single group-valued register. One can also consider *multiregister* experiments, in which we carry out k queries of f , prepare k independent coset states, and then perform a joint measurement on the product state $\rho = \rho^{\otimes k} = \rho \otimes \cdots \otimes \rho$. Note that this measurement does not generally consist of k independent measurements; rather, it is an *entangled* measurement, in which we measure vectors in $\mathbb{C}[G^k]$ along a basis whose basis vectors are not tensor products of k basis vectors in $\mathbb{C}[G]$. For instance, Ip [19] showed that the optimal measurement in the dihedral group is already entangled in the two-register case.

In one sense we already know that such a measurement can succeed. Ettinger, Høyer and Knill [8] showed that the density matrices ρ become nearly orthogonal for distinct subgroups for some $k = O(\log |G|)$. As a consequence, a measurement exists which determines the hidden subgroup with high probability. In [9] they make this result somewhat more constructive by giving an algorithm which solves the HSP by performing a brute-force search through the subgroup lattice of G ; however, for groups of interest such as the symmetric groups, this algorithm takes exponential time. For the dihedral groups in particular, Kuperberg [23] devised a subexponential algorithm, which uses $2^{O(\sqrt{\log n})}$ time and registers, that works by combining two registers at a time and decomposing them into irreducibles.

Regev [30] provided a beautiful kind of worst-case to average-case quantum reduction, by showing that the HSP for the dihedral group D_n can be reduced to uniformly random instances of the Subset Sum problem on \mathbb{Z}_n . Bacon, Childs, and van Dam [2] deepened this connection by determining the *optimal* multiregister measurement for the dihedral group, and showing that it consists of the so-called *pretty good measurement* (PGM); they used this to show a sharp threshold at $k = \log_2 n$ for the number of registers needed to solve the HSP. Moore and Russell [24] generalized their results to some extent, showing that the PGM is optimal

for arbitrary groups G in the single-register case whenever we wish to distinguish the conjugates of some subgroup H from each other, and optimal in the multiregister case whenever (G, H) form a Gel'fand pair.

Whether a similar approach can be taken to the symmetric group S_n is a major open question. In particular, we would like to know whether there is a worst-case to average-case reduction analogous to Regev's, connecting the HSP to some Subset Sum-like problem and whether this would result in new subexponential-time quantum algorithms for Graph Isomorphism. Some recent results show that indeed any such measurement requires a high degree of entanglement: Moore and Russell [26] showed that performing strong Fourier sampling on two registers in S_n requires a superpolynomial number of experiments (specifically, $e^{\Omega(\sqrt{n}/\log n)}$) to distinguish order-2 subgroups $\{1, m\}$ from the identity, or from each other, and conjectured that $\Omega(n \log n)$ registers are necessary. Hallgren, Moore, Rötteler, Russell and Sen [16] proved this conjecture, showing that $\Theta(n \log n)$ registers are necessary and sufficient. Interestingly, the variance over m in the observed probability distribution in the multiregister case has a term for each subset of the registers [25], pointing towards an algorithm that finds a subset with particularly high variance, and thus gives a large amount of information about the hidden subgroup.

Our contribution. In this paper, we consider the special case of the HSP relevant to Graph Isomorphism: namely, where we wish to distinguish the conjugates of some subgroup H from the trivial subgroup, where H has a “missing harmonic” (defined below). We give an explicit k -register measurement in the Fourier basis that distinguishes these two cases. Our approach relies on decomposing the tensor product of the representations observed in a given subset of the registers into a direct sum of irreducibles. Each subset of the registers contributes a small amount of information, so that when $k \geq \log_2 |G|$ the measurement succeeds with constant probability. We hope that this may lead to worst-case to average-case quantum reductions involving generalizations of the Subset Sum problem.

2 Missing Harmonics

We start by preparing independent coset states in k independent G -valued registers, giving the tensor product

$$\rho = \rho^{\otimes k} = \left(\frac{1}{|G|} \sum_{c \in G} |cH\rangle \langle cH| \right)^{\otimes k} = \frac{1}{|G|^k} \sum_{c \in G^k} |cH^k\rangle \langle cH^k|. \quad (1)$$

Note that ρ can also be thought of as a random left coset of the product subgroup $H^k \in G^k$. Note also that ρ is the completely mixed state over $\mathbb{C}[G]^{\otimes k} = \mathbb{C}[G^k]$ if H is the trivial subgroup $\{1\}$.

Here we give an explicit measurement in the Fourier basis which solves an important special case of the HSP, including the case relevant to Graph Isomorphism: namely, given a (non-normal) subgroup $H \subset G$, we wish to distinguish the conjugates of H from the trivial subgroup. Our measurement succeeds with constant probability whenever $k \geq \log_2 |G|$.

Recall that for any representation τ , the average of τ over a subgroup H is a projection operator, which we denote $\tau(H) = (1/|H|) \sum_{h \in H} \tau(h)$. Note that $\tau(H)$ is generally not of full rank, and indeed $\tau(H) = \mathbb{1}_{d_\tau}$ if and only if H is contained in the kernel of τ . Let us say that an irreducible representation η is a *missing harmonic* of H if $\eta(H) = 0$; this is then true for all of H 's conjugates as well. For instance, if G is the dihedral group D_n and $H = \{1, m\}$ where m is one of the n “flips,” then the sign representation π is a missing harmonic. Similarly, if G is the symmetric group S_{2n} where n is odd and H is the order-2 subgroup corresponding to an isomorphic pair of rigid graphs, then the sign representation π , which takes even and odd permutations to $+1$ and -1 respectively, is a missing harmonic.

For simplicity, we focus on the case where H has some missing harmonic η ; the idea is that if we ever observe it, then we know that the hidden subgroup must be trivial rather than a conjugate of H . The following lemma gives some sufficient conditions for H to have a missing harmonic; these are intended as examples, and are by no means exhaustive.

Lemma 1. *If any of the following conditions hold, then H has a missing harmonic:*

1. H is normal and nontrivial.
2. H intersects every coset of some proper normal subgroup $K \triangleleft G$.
3. $G = S_n$ and H is transitive.
4. $|G|/|H| < C$ where $C = \sum_{\tau \in \widehat{G}} d_\tau$.

Proof. 1) Recall that if H is normal then for every $\tau \in \widehat{G}$, either $\tau(H) = \mathbb{1}_{d_\tau}$ or $\tau(H) = 0$. If H is not the trivial subgroup, then the latter must be true for at least one τ .

2) Recall that any irreducible representation of G/K gives an irreducible representation τ of G by composing it with the homomorphism $\phi : G \rightarrow G/K$. Since $\phi(H) = G/K$, we have $\tau(H) = 0$ for any such τ other than the trivial representation. (For instance, in Graph Isomorphism where n is odd, H is transverse to the alternating group A_n and τ is the sign representation.)

3) Let τ be the *standard representation*, corresponding to the Young diagram $(n-1, 1)$. This permutes the n vertices of an $(n-1)$ -dimensional simplex centered at the origin. If H is transitive, then for any i, j exactly $1/n$ of the elements of H take vertex i to vertex j , and so the average $\tau(H) = 0$ is zero.

4) Recall that the *regular representation* Reg , namely $\mathbb{C}[G]$ under left multiplication by G , consists of d_τ copies of each $\tau \in \widehat{G}$. It is a simple exercise to show that $\text{rk } \text{Reg}(H)$ is the index $|G|/|H|$, and since

$$\text{rk } \text{Reg}(H) = \sum_{\tau \in \widehat{G}} d_\tau \text{rk } \tau(H)$$

we have $\text{rk } \text{Reg}(H) \geq C$, a contradiction, if $\tau(H) \neq 0$ for all $\tau \in \widehat{G}$. □

3 Decomposing Subsets of the Registers

The state ρ is a density matrix defined on the Hilbert space $\mathbb{C}[G^k] = \mathbb{C}[G]^{\otimes k}$. Since it is completely mixed over left cosets of H^k , it commutes with left multiplication in G^k . It follows from Schur's lemma [23, 25, 26] that ρ is block-diagonal in the Fourier basis, where each block corresponds to one of the irreducible representations of G^k . These are tensor products of irreducible representations of G , $\sigma = \sigma_1 \otimes \cdots \otimes \sigma_k$. To put it differently, the optimal measurement is consistent with first performing weak Fourier sampling on each of the k registers, observing the representation names $\sigma_1, \dots, \sigma_k$.

The question is how to refine this measurement further, decomposing σ into smaller subspaces. (We abuse notation by identifying subspaces with the name of the representation that acts on them.) Happily, there is a natural way to do this that respects the structure of G : specifically, we treat σ as a representation of G (rather than of G^k) by restricting to the *diagonal action*, where the element $g \in G$ acts by $\sigma(g) = \sigma_1(g) \otimes \cdots \otimes \sigma_k(g)$. We can then further decompose σ into irreducible representations $\tau \in \widehat{G}$ under this action. If we observe a missing harmonic η under this decomposition, we know that the hidden subgroup is trivial rather than being a conjugate of H . Unfortunately, in all cases of interest η is very low-dimensional (indeed, one-dimensional), and so the chances of observing η are exponentially small even if the hidden subgroup is trivial. Thus this direct approach does not work.

Instead, we focus on some subset $I \subseteq [k]$ of the registers. First, we can decompose σ into the tensor product of the registers inside and outside I , $\sigma = (\bigotimes_{i \in I} \sigma_i) \otimes (\bigotimes_{i \notin I} \sigma_i)$. Now consider the decomposition of the registers in I into irreducible representations of G under the diagonal action, in which we right-multiply² every register in I by g and leave the other registers fixed. We write $\bigotimes_{i \in I} \sigma_i = \tau_1 \oplus \cdots \oplus \tau_\ell$. Fixing our missing harmonic η , for each nonempty I this gives us a subspace

$$W_{\eta, \sigma}^I = \left(\bigoplus_{i: \tau_i \cong \eta} \tau_i \right) \otimes \left(\bigotimes_{i \notin I} \sigma_i \right) ,$$

²We use right multiplication because left cosets of H are invariant under right multiplication by H .

and we define $\Pi_{\eta,\sigma}^I$ as the projection operator which projects onto this subspace. That is, $\Pi_{\eta,\sigma}^I$ projects the registers in I into irreducible subspaces isomorphic to η , and leaves the other registers fixed. The following lemma shows that if η is a missing harmonic for H , then each of these projection operators annihilates ρ .

Lemma 2. *Suppose that $\eta(H) = 0$. Then for all nonempty $I \subseteq [k]$ and all σ , $\Pi_{\eta,\sigma}^I \rho = 0$.*

Proof. The state ρ is symmetric under right multiplication by any $\mathbf{h} \in H^k$. In particular, it is symmetric under right diagonal multiplication by any $h \in H$ on the registers in I . Let R_H^I be the operator which symmetrizes over this action: that is, the average over all $h \in H$ of the unitary operator that right-multiplies by \mathbf{h} where $h_i = h$ for $i \in I$ and $h_i = 1$ for $i \notin I$. Then

$$\rho = R_H^I \rho (R_H^I)^\dagger \quad \text{and} \quad \Pi_{\eta,\sigma}^I R_H^I = 0$$

and so

$$\Pi_{\eta,\sigma}^I \rho = (\Pi_{\eta,\sigma}^I R_H^I) \rho (R_H^I)^\dagger = 0 .$$

□

Now we patch these operators together to form our measurement. Let

$$W_{\eta,\sigma} = \text{span}_{I \subseteq [k]} W_{\eta,\sigma}^I$$

be the span of all these subspaces, and let $\Pi_{\eta,\sigma}$ be the projection operator onto $W_{\eta,\sigma}$. By Lemma 2, we know that $\Pi_{\eta,\sigma} \rho = 0$ whenever η is a missing harmonic for the hidden subgroup. Thus we can distinguish the conjugates of H from the trivial subgroup with a measurement operator that reports “trivial” if it observes the subspace $W_{\eta,\sigma}$, and “don’t know” if it observes the perpendicular subspace $W_{\eta,\sigma}^\perp$. Since ρ is completely mixed if the hidden subgroup is trivial, the probability that our operator reports “trivial” in that case is $\dim W_{\eta,\sigma} / d_\sigma$. We wish to show that if $k \geq \log_2 |G|$, the expectation over σ of this fraction is at least $1/2$, so that our measurement distinguishes the trivial subgroup from conjugates of H with constant probability.

To calculate this expectation, it is convenient to work in the entire Hilbert space $\mathbb{C}[G^k]$ of the k registers, rather than conditioning on having observed the representation names σ . Recall that the action of G on $\mathbb{C}[G]$ under (right) group multiplication yields the regular representation Reg , and that Reg contains d_σ copies of each $\sigma \in \widehat{G}$. It follows that the fraction of $\mathbb{C}[G]$, dimensionwise, consisting of copies of σ is $d_\sigma^2 / |G|^k$. This fraction is also the probability that we observe the representation name σ in a given register when we perform weak Fourier sampling on the completely mixed state, and is called the *Plancherel distribution* $\text{Planch}(\sigma)$. Similarly, $\mathbb{C}[G^k]$ can be thought of as the regular representation of G^k , in which case it contains $d_\sigma = \prod_i d_{\sigma_i}$ copies of each σ , giving the Plancherel distribution $\text{Planch}(\sigma) = \prod_i \text{Planch}(\sigma_i)$. Thus we have

$$\text{Exp}_\sigma \frac{\dim W_{\eta,\sigma}}{d_\sigma} = \sum_\sigma \text{Planch}(\sigma) \frac{\dim W_{\eta,\sigma}}{d_\sigma} = \frac{\sum_\sigma d_\sigma W_{\eta,\sigma}}{|G|^k} = \frac{\dim W_\eta}{|G|^k} .$$

In other words, the expected dimensionwise fraction of $W_{\eta,\sigma}$ in σ is the total dimensionwise fraction of W_η in all of $\mathbb{C}[G^k]$, where

$$W_\eta = \text{span}_\sigma W_{\eta,\sigma} .$$

We can also write

$$W_\eta = \text{span}_{I \subseteq [k]} W_\eta^I$$

where

$$W_\eta^I = \text{span}_\sigma W_{\eta,\sigma}^I$$

is the subspace of $\mathbb{C}[G^k]$ spanned by vectors for which, if we decompose the registers in I into G -irreducibles, we observe the representation name η , regardless of what G^k -representation σ they lie in.

We wish to lower bound the fraction of $\mathbb{C}[G^k]$ consisting of W_η . First, we ask how much of $\mathbb{C}[G^k]$ consists of each W_η^I . Recall that for any representation ϕ , $\phi \otimes \text{Reg}$ consists of the direct sum of d_ϕ copies of Reg . In particular, $\text{Reg}^{\otimes \ell}$ contains $|G|^{\ell-1} d_\sigma$ copies of each $\sigma \in \widehat{G}$. Thus for any I , if $|I| = \ell$ we have

$$W_\eta^I \cong |G|^{\ell-1} d_\eta \eta \otimes \text{Reg}^{\otimes(k-\ell)}$$

and so

$$\frac{\dim W_\eta^I}{|G|^k} = \frac{|G|^{k-1} d_\eta^2}{|G|^k} = \frac{d_\eta^2}{|G|} . \quad (2)$$

In other words, each W_η^I occupies the same fraction of $\mathbb{C}[G^k]$ as η occupies of $\mathbb{C}[G]$, namely the Plancherel distribution. This is just $1/|G|$ in the cases we care about, but then again there are 2^k subsets I . Our hope is that when $k \sim \log_2 |G|$, then, the span of all the W_η^I occupies a large fraction of the Hilbert space.

Indeed, if the subspaces W_η^I for different I were orthogonal, their dimensions would simply add, giving $\dim W_\eta = (2^k - 1) \dim W_\eta^I$; however, it is easy to see (even for $G = \mathbb{Z}_2$ and $k = 2$) that this is not the case. Instead, it turns out that the subspaces W_η^I have a remarkable statistical property akin to pairwise independence, which we describe in the next section.

4 Independent Subspaces

We say that two subspaces W_1, W_2 are *independent* if the expected squared projection of a random vector $v \in W_1$ into W_2 is just what it would be if v were a random vector in the entire space, i.e., the dimensionwise fraction of that space occupied by W_2 . This is a kind of statistical independence between the events that we observe W_1 and W_2 (although if their projection operators do not commute, we cannot consider these simultaneously as quantum observables!) Formally:

Definition 1. Let V be a vector space. Let W_1, W_2 be subspaces with projection operators Π_1, Π_2 . Let w be chosen uniformly at random from the vectors in W_1 with norm 1. Then W_1 and W_2 are independent if

$$\text{Exp}_w |\Pi_2 w|^2 = \frac{\dim W_2}{\dim V} .$$

Equivalently, if v is chosen uniformly at random from the vectors in V with norm 1,

$$\text{Exp}_v |\Pi_1 \Pi_2 v|^2 = \frac{\dim W_1}{\dim V} \frac{\dim W_2}{\dim V} .$$

A family of subspaces W_1, \dots, W_m is independent if W_i and W_j are independent for any distinct i, j .

Note that this definition remains the same if we choose w or v uniformly from an orthonormal basis for W_1 or V respectively, rather than from the sphere of radius 1. Indeed, since $\text{Exp}_v |\Pi_1 \Pi_2 v|^2 = \text{tr } \Pi_1 \Pi_2 / \dim V$, a more compact definition is the following:

$$\text{tr } \Pi_1 \Pi_2 = \frac{\text{tr } \Pi_1 \text{tr } \Pi_2}{\dim V} = \frac{\dim W_1 \dim W_2}{\dim V} . \quad (3)$$

Note, however, that $\Pi_1 \Pi_2$ is not a projection operator unless Π_1 and Π_2 commute.

Lemma 3. Let $I, J \subseteq [k]$ be distinct and nonempty. Then W_η^I and W_η^J are independent.

Proof. We use the fact that for any representation U and an irreducible representation τ , the projection operator onto the *isotypic subspace* corresponding to τ —that is, the span of all the copies of τ in U —is

$$\Pi_\tau = \frac{d_\tau}{|G|} \sum_{g \in G} \chi_\tau(g)^* U(g) .$$

In particular, if Π_η^I projects onto W_η^I , we have

$$\Pi_\eta^I = \frac{d_\eta}{|G|} \sum_{g \in G} \chi_\eta(g)^* \text{Reg}(g)^{\otimes I} \otimes \mathbb{1}^{\otimes ([k] \setminus I)}$$

and

$$\Pi_\eta^I \Pi_\eta^J = \left(\frac{d_\eta}{|G|} \right)^2 \sum_{g, g' \in G} \chi_\eta(g)^* \chi_\eta(g')^* \text{Reg}(g)^{\otimes (I \setminus J)} \otimes \text{Reg}(gg')^{\otimes (I \cap J)} \otimes \text{Reg}(g')^{\otimes (J \setminus I)} \otimes \mathbb{1}^{\otimes ([k] \setminus (I \cup J))} ,$$

where the notation $A^{\otimes I}$ denotes the tensor product $\otimes_{i \in I} A$. Taking traces, since $\chi_{\text{Reg}}(g) = |G|$ for $g = 1$ and 0 for $g \neq 1$, whenever $I \neq J$ the summand is zero unless both $g = 1$ and $g' = 1$. In this case the summand is $d_\eta^2 |G|^k$, so

$$\text{tr } \Pi_\eta^I \Pi_\eta^J = \left(\frac{d_\eta^2}{|G|} \right)^2 |G|^k .$$

Since $\text{tr } \Pi_\eta^I = \text{tr } \Pi_\eta^J = (d_\eta^2 / |G|) |G|^k$ by (2), we see that Π_η^I and Π_η^J satisfy (3) and are independent. \square

Finally, we lower bound the dimension of the span of a independent family of subspaces with the following lemma, and show $\dim W_\eta / |G|^k \geq 1/2$ whenever $k \geq \log_2 |G|$.

Lemma 4. *Let V have dimension D , and let $W_1, \dots, W_m \subset V$ be a independent family of subspaces of dimension d , and let $W = \text{span}_i W_i$. Then*

$$\frac{\dim W}{D} \geq 1 - \frac{1}{1 + md/(D - d)} .$$

Proof. Let Π_i project onto W_i for each $1 \leq i \leq m$, and consider the operator $M = \sum_{i=1}^m \Pi_i$. Since M is positive and symmetric, it can be diagonalized, and has nonzero eigenvalues $\lambda_1, \dots, \lambda_t > 0$ where its rank is $t = \dim W$. Its trace is

$$\sum_{\ell=1}^t \lambda_\ell = md . \quad (4)$$

Then using (3), the Frobenius norm of M is

$$\|M\|^2 = \sum_{\ell=1}^t \lambda_\ell^2 = \text{tr } M^\dagger M = \sum_{i,j} \text{tr } \Pi_i \Pi_j = md + \sum_{i \neq j} \text{tr } \Pi_i \Pi_j = md + m(m-1) \frac{d^2}{D} . \quad (5)$$

On the other hand, by Cauchy-Schwartz we have

$$\left(\sum_{\ell=1}^t \lambda_\ell \right)^2 \leq t \sum_{\ell=1}^t \lambda_\ell^2 \quad (6)$$

Combining (4), (5) and (6) gives

$$md \leq t \left(1 + (m-1) \frac{d}{D} \right)$$

and so

$$\frac{t}{D} \geq \frac{md}{D + (m-1)d} = 1 - \frac{1}{1 + md/(D-d)} .$$

\square

Applying this to the independent family $\{W_\eta^I \mid I \subseteq [k], I \neq \emptyset\}$ gives the following corollary.

Corollary 1. *For any $k \geq \log_2 |G|$, we have $\dim W_\eta / |G|^k \geq \frac{1}{2}$.*

Proof. We have $V = \mathbb{C}[G^k]$, $D = |G|^k$, $m = 2^k - 1 \geq |G| - 1$, and $d/D = d_\eta^2 / |G| \geq 1/|G|$. Thus $md/(D-d) \geq 1$, and Lemma 4 completes the proof. \square

5 The Representation Kickback Trick

In this section we show how to efficiently carry out the von Neumann measurement associated with the subspace W_η^I for a fixed subset I of the registers. Of course, this does not tell us how to efficiently carry out the measurement associated with their span W_η .

It suffices to consider the space $V = \bigotimes_{i \in I} \sigma_i$, decompose $V = \bigoplus_{\tau \in \widehat{G}} a_\tau \tau$ into irreducible representations of G , and implement the measurement associated with the projection operator Π_η that projects onto the space spanned by the a_η copies of η in this direct sum above. Our approach is essentially the same as the “summand extraction” of Kuperberg [23].

To carry out this measurement, we introduce a new G -valued control register, in which we initially prepare $|G\rangle$, the uniform superposition over G . Treating our state now as an element of $\mathbb{C}[G] \otimes V$, we apply the *controlled G -action* operator: $\mathcal{M} : |g\rangle \otimes |\phi\rangle \mapsto |g\rangle \otimes \sigma_I(g^{-1})|\phi\rangle$, where $\sigma_I(g^{-1})$ is the unitary operator $\bigotimes_{i \in I} \sigma_i(g^{-1})$. Finally, we compute the quantum Fourier transform on the control register and carry out the measurement (on the control register only) corresponding to the operators Π_η and $\mathbb{1} - \Pi_\eta$, where $\Pi_\eta : \mathbb{C}[G] \rightarrow \mathbb{C}[G]$ is the operator that projects onto the irreducible subspaces isomorphic to η .

To see why this works, we return our attention to $\mathbb{C}[G] \otimes V$. Consider the following two G -actions on this space: let $D_h : |g\rangle \otimes |\phi\rangle \mapsto |hg\rangle \otimes \sigma_I(h)|\phi\rangle$ apply the group action to *both* the control register and V , and let $L_h : |g\rangle \otimes |\phi\rangle \mapsto |hg\rangle \otimes |\phi\rangle$ apply the group action only on the control register. Then observe that $\mathcal{M} \circ D_h = L_h \circ \mathcal{M}$: evidently, any subspace of $\mathbb{C}[G] \otimes V$ that is invariant under D_h is (unitarily) transformed by \mathcal{M} to a subspace that is invariant under L_h . Observe now that if $W \subset V$ is an invariant subspace of V (under $\sigma_I(h)$) that is isomorphic to η , then $|G\rangle \otimes W$ is isomorphic to η under the action D_h , as $|G\rangle$ is invariant under right multiplication by G . Thus such a space is carried to a L_h -invariant space by \mathcal{M} , still isomorphic to η .

6 Discussion

We have suggested here a general framework for solving cases of the Hidden Subgroup Problem similar to that relevant to Graph Isomorphism. Each subset of the registers contributes a small amount of information, and the span of all their associated subspaces distinguishes the trivial subgroup from the conjugates of a non-trivial subgroup with a missing harmonic. Of course, even though we can project into each of these subspaces efficiently, it is far from clear how to product into their span. However, we might hope, through some partial measurement, to find an *informative subset*: that is, a subset I for which a large fraction of the state lies in W_η^I if the hidden subgroup is trivial.

One approach to finding such a subset is a “sieve,” in which we combine states in pairs, project their tensor products into irreducible representations, and try to work our way down towards a missing harmonic. This builds a tree of tensor products with a missing harmonic at its root, and the chosen subset of registers corresponds to the leaves of the tree. This idea is not original with us: if we consider the special case of the HSP in the dihedral group D_n where we wish to distinguish the trivial subgroup from the n conjugate subgroups $H = \{1, m\}$ where m is a “flip,” this is exactly what Kuperberg’s algorithm does. Namely, it combines two-dimensional representations, decomposing them according to $\sigma_i \otimes \sigma_j \cong \sigma_{i+j} \oplus \sigma_{i-j}$, until we reach $\sigma_0 \cong \mathbb{1} \oplus \pi$. We then perform a measurement inside σ_0 , and if we ever observe the sign representation π we know that the hidden subgroup is trivial.

Using a somewhat different type of sieve, Alagic, Moore and Russell [1] recently obtained a subexponential-time algorithm for the HSP in groups of the form G^n for finite G . Even though these groups have a simple structure, they are similar to S_n in that most of their irreducible representations are exponentially high-dimensional, and it was shown in [16] that solving their HSP requires entangled measurements over $\Omega(n)$ registers. Whether a similar sieve can work for S_n , providing an efficient (even subexponential) quantum algorithm for Graph Isomorphism remains an open question.

Acknowledgments

We are grateful to Dorit Aharonov, Andrew Childs, Tracy Conrad, Gabor Ivanyos, Greg Kuperberg, Sally Milius, Rosemary Moore, Martin Rötteler, Leonard Schulman, Pranab Sen, Douglas Strain and Umesh Vazirani for helpful discussions. We also thank the organizers of the Banff Workshop on Quantum Computation in September 2004, where some of this work was done. We gratefully acknowledge the support of the NSF through grants EIA-0218563, CCR-0220070, and CCF-0524613, and ARO contract W911NF-04-R-0009.

References

- [1] Gorjan Alagic, Cristopher Moore, and Alexander Russell. Subexponential-time algorithms for hidden subgroup problems over product groups. Preprint, quant-ph/0603251 (2006).
- [2] David Bacon, Andrew Childs, and Wim van Dam. Optimal measurements for the dihedral hidden subgroup problem. *Chicago Journal of Theoretical Computer Science*, to appear.
- [3] David Bacon, Andrew Childs, and Wim van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. *Proc. 46th Foundations of Computer Science*, pages 469–478, 2005.
- [4] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. *Proc. 29th ACM Symp. on Theory of Computing*, pages 48–53, 1997.
- [5] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory (preliminary abstract). *Proc. 25th ACM Symp. on Theory of Computing*, pages 11–20, 1993.
- [6] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *Proc. 14th ACM-SIAM Symp. on Discrete Algorithms*, pages 489–498, 2003.
- [7] Mark Ettinger and Peter Høyer. On quantum algorithms for noncommutative hidden subgroups. Preprint, quant-ph/9807029 (1998).
- [8] Mark Ettinger and Peter Høyer and Emmanuel Knill. Hidden subgroup states are almost orthogonal. Preprint, quant-ph/9901034.
- [9] Mark Ettinger and Peter Høyer and Emmanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, to appear.
- [10] Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. *Proc. 35th ACM Symp. on Theory of Computing*, 2003.
- [11] William Fulton and Joe Harris. *Representation Theory: A First Course*. Number 129 in Graduate Texts in Mathematics. Springer-Verlag, 1991.
- [12] Michelangelo Grigni, Leonard J. Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Proc. 33rd ACM Symp. on Theory of Computing*, pages 68–74, 2001.
- [13] Lisa Hales and Sean Hallgren. Quantum fourier sampling simplified. *Proc. 31st ACM Symp. on Theory of Computing*, 1999.
- [14] Lisa Hales and Sean Hallgren. An improved quantum Fourier transform algorithm and applications. *Proc. 41st Symp. on Foundations of Computer Science*, 2000.
- [15] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. *Proc. 32nd ACM Symp. on Theory of Computing*, pages 627–635, 2000.

- [16] Sean Hallgren, Christopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen, Limitations of quantum coset states for Graph Isomorphism. *Proc. 38th ACM Symp. on Theory of Computing* (STOC 2006), to appear.
- [17] Peter Høyer. Efficient quantum transforms. Preprint, quant-ph/9702028 (1997).
- [18] Yoshifumi Inui and François Le Gall. An efficient algorithm for the hidden subgroup problem over a class of semi-direct product groups. Proc. EQIS 2004.
- [19] Lawrence Ip. Shor's algorithm is optimal. Preprint, 2004.
- [20] Gábor Ivanyos, Frédéric Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. *Int. J. Found. Comput. Sci.* 14(5): 723–740, 2003.
- [21] Richard Jozsa. Quantum factoring, discrete logarithms and the hidden subgroup problem. Preprint, quant-ph/0012084 (2000).
- [22] Julia Kempe and Aner Shalev, The hidden subgroup problem and permutation group theory. *Proc. 16th ACM-SIAM Symp. on Discrete Algorithms*, pages 1118–1125, 2005.
- [23] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. Preprint, quant-ph/0302112 (2003).
- [24] Christopher Moore and Alexander Russell. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets. Preprint, quant-ph/0501177
- [25] Christopher Moore, Alexander Russell, and Leonard Schulman. The symmetric group defies strong Fourier sampling. *Proc. 46th Foundations of Computer Science* (FOCS 2005), 479–490.
- [26] Christopher Moore and Alexander Russell. The symmetric group defies strong Fourier sampling: Part II. Preprint, quant-ph-0501066.
- [27] Christopher Moore, Daniel Rockmore, and Alexander Russell. Generic quantum Fourier transforms. *Proc. 15th ACM-SIAM Symp. on Discrete Algorithms*, pages 778–787, 2004.
- [28] Christopher Moore, Daniel Rockmore, Alexander Russell, and Leonard J. Schulman, The value of basis selection in Fourier sampling: hidden subgroup problems for affine groups. *Proc. 15th ACM-SIAM Symp. on Discrete Algorithms*, pages 1113–1122, 2004.
- [29] Jaikumar Radhakrishnan, Martin Rötteler, and Pranab Sen. On the Power of Random Bases in Fourier Sampling: Hidden Subgroup Problem in the Heisenberg Groups. *Proc. 32nd Intl. Colloquium on Automata, Languages, and Programming*, pages 1399–1411, 2005.
- [30] Oded Regev, Quantum computation and lattice problems. *Proc. 43rd Symp. on Foundations of Computer Science*, pages 520–530, 2002.
- [31] Martin Rötteler and Thomas Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. Preprint, quant-ph/9812070 (1998).
- [32] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Number 42 in Graduate Texts in Mathematics. Springer-Verlag, 1977.
- [33] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [34] Daniel R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.